

ATLAS Data Challenges

15th November 2001

Introduction

About two years ago ATLAS computing planned a first series of Data Challenges in order to validate its Computing Model, its software, its data model, and to ensure the correctness of the technical choices to be made. Since then, and particularly in the context of the CERN Review of LHC Computing, the scope and goals of these data challenges has evolved. It is now clear that these Data Challenges should be executed at the prototype Tier centres which will be built in the Phase1 of the LHC Computing Grid Project. The goals of these data challenges are given below. It is understood that they should be of an increasing complexity and will use as much as possible the Grid middleware being developed in the context of several Grid projects like EU Data Grid or GridPP. The results of these data challenges will be used as input for a Computing Technical Design Report due by the end of 2003 and for preparing a Memorandum of Understanding in due time.

The Data Challenges in Brief

We are considering here the first three data challenges (DC) that will be run from now to the end of 2003. For all of them it is intended to have physics content in order to bring the physicists community into the exercise and for a more sensible validation of the software. For DC1, in 2002, a major goal is to provide simulated data to the High Level Trigger (HLT) community that has to prepare its own TDR by the end of the year. The aim is to generate, reconstruct and analyse the number of Monte Carlo event listed over a 2-3 week period making use of the prototype centres and Grid-enabled software. Data Challenges will of course continue beyond Data Challenge 2, probably annually, but detailed discussion of these has not yet begun.

- Data Challenge 0 runs from November-December 2001 and is to test the continuity of the code chain. Only modest samples will be generated, and essentially all in “traditional” sequential file format. In order to test the database technology, today Objectivity, the simulated data will be converted to Objectivity before the reconstruction step. In addition it is also intended to convert part of the data simulated with the old software, for the Physics TDR in 1998-99, to the new technology.
- Data Challenge 1 runs in the first half of 2002, and will be divided in two phases. In the first phase we will generate several data sets, some as large as 10^7 events, for high-level trigger studies and for physics analysis. The second phase will also produce some physics samples but will be more oriented to the testing of the new software, including Geant4, the new event data model and the evaluation of two database technologies, Objectivity and Root-I/O. As stated above, the aim is to be able to produce a sample of around 10^7 events over 2-3 weeks. Based on our current knowledge the volume of data to be generated should be of the order of 20-30 TB. DC1 will run on sites worldwide, and will make as much use of Grid tools as time permits.
- Data Challenge 2 runs for the first half of 2003. The scope will depend on what has been and has not been achieved in DC1, but the main goal will be to have the full new software in place. We will generate several samples of 10^8 events, mainly in OO-databases, and with large-scale physics analysis using Grid tools.

All the Data Challenges will be run on Linux systems operating at close to the CERN certified level, and with the compilers distributed with the code if not already installed locally in the correct version.

The Goals of the ATLAS Data Challenges

Data Challenge 0:

DC0 comprises a 'continuity' test through the software chain. The aim is primarily to check the state of readiness for Data Challenge 1. We intend to simulate and reconstruct samples of the order of 100k "Z+jet" events, or similar. This is the right scale, but the number of events is less important than checking that the software works, the primary goal is to ensure that the full software chain is ready. Issues to be checked include:

- Geant3 simulation on PC farm
- 'pile-up' handling
- reconstruction running
- data must be written/read to/from the database

In addition few samples of single particle data (electrons, muons, photons,) will be generated to finalise the study of the impact of the new detector layout.

In parallel, again in view of DC1, we intend to check that the infrastructure to use Objectivity (writing to and reading from) is fully working and to do so we will convert in a systematic way the existing (part of) Physics TDR data generated in 1997-98 and will run the reconstruction in the Athena framework.

DC0 will largely be run at a single centre, probably CERN; however, the major centres intending to participate in DC1 will be encouraged to perform some generations and to test the software. Note that, assuming a typical PC can simulate approximatively 20 events an hour, then a farm of 100 PCs can generate 100k events in a few days. So the hardware (CPU and storage) should not be a big issue. Of course, we should count on trying to generate the sample 3-4 times during the 6-week period: we are bound to get it wrong the first few times!

Data Challenge 1:

The scope of DC1 increases significantly beyond DC0.

- several samples of up to 10^7 events should be produced
- the production should involve CERN and outside-CERN sites. This will in fact be necessary in order to have enough cpu power and data storage.

As a goal we should aim to have periods when we run on several hundred PCs, worldwide. (1000 PCs can simulate approximatively 10^7 events in 10-20 days.)

The aims of this Data Challenge are several:

1. provide a sample of order 10^7 di-jet events each for High Level Trigger (HLT) studies:
 - o to improve on previous statistics by factor 10;
 - o to study performance of Athena and algorithms for use in HLT;
 The HLT Technical Design Report is due for the end of 2002.
 Pile-up at low (2×10^{33}) and high luminosity (10^{34}) should be included. The events will have to pass through the ATLFAST chain as well.
2. Try out running 'reconstruction and 'analysis' on a large scale. (This is after all what has to be done on 'real' data). We expect we will learn a lot from our first attempts to do this, including whether our data model, I/O, etc. are 'bottle-necks'. See also 5) below.
3. provide samples of events for physics studies. The data generated would be mainly 'Standard Model' backgrounds. We hope to start some detailed checking of Geant3 versus Geant4 towards the end of this DC. (The HLT sample will be generated with G3). Hence we have to understand HOW to do this G3/G4 comparison: issue of 'same' geometry arises.
4. as the production will take place at several sites, we will also be going to gain experience of a 'distributed' model in practice. We would try to make use of a certain amount of GRID tools, as available.
5. use some of the produced data to carry out evaluation of dBase technologies: this means processing N events with the different technologies.

As mentioned above, the Data Challenge has two phases. The first one will be oriented to the production of data to fulfil the request of the high-level trigger community. The second one will be more oriented to the testing of the new software, including Geant4, the new event data model and the evaluation of two database technologies, Objectivity and Root-I/O. It is expected that a significant update of the software will happen between the two phases.

Data Challenge 2:

This Data Challenge should be done between spring and Autumn 2003.

The details of DC2 will depend on what has and has not been achieved in DC0/1. The goal at this stage includes:

- use 'Test-Bed' which will be built in the context the Phase 1 of the "LHC Computing Grid Project" with the scale at a sample of the order of 10^8 events. We would like to test out a system with "complexity" at something like ~50% of 2006-2007 system;
- Geant4 simulation should play a major role in this DC;
- Physics samples generated could have 'hidden' new physics;
- calibration and alignment procedures should be tested;
- extensive use of the GRID middleware;
- open question: how far back up the trigger chain do we need/want to go, at least from some fraction of the DC.

To summarize:

Data Challenge	Size	Date
DC0	10^5 events	November - December 2001
DC1	10^7 events	February – July 2002
DC2	10^8 events	March – September 2003

An estimation of the resources needed for DC1 have been elaborated on the basis of the previous large-scale jet production performed in 1997. Since it corresponds to specific physics channels for the purpose of the High Level Trigger Technical Design Report we name it DC1-HLT. Details on the way the production was made, like type of generated events, selection criteria can be found in the ATLAS Physics note ATL-PHYS-97-102 (30 July 1997).

	Number of events	Time per event SI95 sec	Total time SI95-sec	Total time SI95-hours
Simulation	10^7	3000	3×10^{10}	10^7
Reconstruction	10^7	640	6.4×10^9	2×10^6

Table 2: DC1 CPU resources

	Number of events	Event size MB	Total size GB	Total size TB
Simulation	10^7	2	20000	20
Reconstruction	10^7	0.5	5000	5

Table 3: DC1 Volume of data

L	Number of events	Event size MB	Total size GB	Total size TB
2×10^{33}	1.5×10^6	(1) 2.6	4000	4
		(2) 4.7	7000	7
10^{34}	1.5×10^6	(1) 6.5	10000	10
		(2) 17.5	26000	26

Table 4: DC1 volume of data with pile-up for 2 luminosities

The numbers are given for the events to be simulated for the HLT trigger. It assumes a filtering after the detector simulation in which ~14% of the events are kept.

- (1) One assumes that only Monte Carlo ‘digits’ are written to the output stream.
- (2) One assumes that both Monte Carlo ‘hits’ and ‘digits’ are written to the output stream.

ATLAS Application Descriptions

ATLAS proposes to develop their distributed Monte Carlo production system as a short term use case for developing, in an incremental fashion, essential components of the Datagrid software. The ATLAS MC production system has the following components:

- 1) Job submission system
- 2) Job scripts running the executables
- 3) Temporary and mass storage of output datasets
- 4) Copying of part of the data onto tape at CERN
- 5) Bookkeeping (meta) database/catalogue.

As of today, ATLAS uses two different frameworks. Atlsim, which has been developed for the early studies of the detector and the Athena/Gaudi framework, developed in collaboration with the LHCb experiment.

The event generation can use several event generators (e.g. Pythia, Herwig, Isajet, etc.) and can run either in the old Fortran Atlsim framework or in the new Athena framework. The output can be written in Objectivity using HepMc and in zebra.

The fast simulation, Atlfast, has been moved to OO. It runs in the Athena framework, can read HepMC format data from Objectivity and produce either Objectivity data or classical “n-tuples”.

The current detector simulation code called DICE is Fortran based, uses Geant 3.21 to track the events through the detector and runs in the Atlsim framework. Events are written out in the form of Zebra banks. New detector simulation based on Geant4 is being developed, we expect that part of the data simulated in DC1 will be done with this new simulation and that it will play a major role in DC2.

Most of the reconstruction programs have now moved to OO, even if some packages are still in Fortran. The new reconstruction uses the Athena framework and in the current situation input data can come from zebra or from Objectivity.

Essential components required for ATLAS Monte Carlo production are the associated bookkeeping and meta-data services. The bookkeeping system needs to be developed to allow replications to occur at remote sites and, very importantly, synchronisation tools to ensure the integrity of a distributed database system. This system should deal not only with the input and output of simulated data but equally vital is the description of the detector geometry. The Monte Carlo production tools will need to interface to the Grid services for data location and transfer mechanisms and Grid enabled I/O to ensure this will occur. The access to mass storage should be using a uniform environment across the Grid via the Grid uniform naming scheme.

ATLAS Generic Use Cases

The necessary software to perform the tasks outlined below is assumed to be available at all Atlas enabled remote sites, as it is included in the Atlas Software Distribution kit. The use cases for a standard software distribution on the Grid are given elsewhere. We note that the first use case is particularly suited to realisation in the first year (by the end of 2002). Indeed, the following items are assumed to be available by March 2002:

- Job submission via a rough resource broker that takes into account the presence of the input data (specified by whole datasets) for choosing a site
- Interface to a metadata bookkeeping service for the job info.
- Interface to a replica catalogue (data metadata) and to a prototype data management (Magda)
- Access to a job monitoring system
 - For some sites, the storage of data in CASTOR or HPSS

By DC2 in March 2003, a preliminary version of an interface to the GRID services is assumed to be in production. This should allow event selection of more complex specification such as physics event classes decided by the event filter. At that stage, comprehensive storage brokering that will allow the automated storage and cataloging of large volumes of data. The ability to deploy the ATLAS code and run-time environment by a “push” mechanism should be in place, and services should allow one to test if the environment is properly in place before the job is dispatched to a given remote site.